

AD-769 673

NATURAL LANGUAGE INFERENCE

Yorick Wilks

Stanford University

Prepared for:

Advanced Research Projects Agency
Department of Defense

August 1973

DISTRIBUTED BY:



National Technical Information Service
U. S. DEPARTMENT OF COMMERCE
5285 Port Royal Road, Springfield Va. 22151

DISCLAIMER NOTICE

**THIS DOCUMENT IS THE BEST
QUALITY AVAILABLE.**

**COPY FURNISHED CONTAINED
A SIGNIFICANT NUMBER OF
PAGES WHICH DO NOT
REPRODUCE LEGIBLY.**

AD 769673

STANFORD RESEARCH INSTITUTE INTELLIGENCE LABORATORY
MEMO NO. AI-211

AUGUST 1973

COMPUTER SCIENCE DEPARTMENT
REPORT NO. CS-383

NATURAL LANGUAGE INFERENCE

by

Yorick Wilks

ABSTRACT: The paper describes the way in which a Preference Semantics system for natural language analysis and generation tackles a difficult class of anaphoric inference problems (finding the correct referent for an English pronoun in context); those requiring either analytic/conceptual knowledge of a complex sort, or requiring weak deductive knowledge of the course of events in the real world. The method employed converts all available knowledge to a canonical template form and endeavors to create chains of non-deductive inferences from the unknowns to the possible referents. Its method of selecting among possible chains of inferences is consistent with the overall principle of "semantic preference" used to set up the original meaning representation, of which these anaphoric inference procedures are a manipulation.

This paper owes much to discussions with Annette Herakovits, Bill Simpson, Ken Colby, Bruce Anderson and Horace Enea. The mistakes are all mine of course.

The views and conclusions contained in this document are those of the author and should not be interpreted as representing necessarily the official policies, either expressed or implied, of the Advanced Research Projects Agency or the U. S. Government.

This research was supported by the Advanced Research Projects Agency, Department of Defense (SD 183), USA.

Reproduced in the USA. Available from the National Technical Information Service, Springfield, Virginia, 22161.

20

INTRODUCTION

The paper describes inferential manipulations of a representation of the meaning structure of natural language. It differs from previous descriptions (1, 2, 3, 4) of this semantics-based system, which have concentrated on the representation itself, and above all on the procedures by which the representation is produced from input sentences and paragraphs in English. In this paper I assume that structure of representation, except for a brief recapitulation, and concentrate on operations upon it for the solution of a class of difficult anaphora problems.

The system described is part of a running system for understanding and translating natural language on the PDP6/10 at the Stanford A. I. Laboratory, programmed in MLISP and LISPL. I shall not in anyway stress the translation-into-French aspect of the work, but its presence provides a continual empirical check of the adequacy of the "reference" and "understanding" described here.

The rather emphasis on the construction of the linguistic base is, I believe, fully justified. The present system is, to my knowledge, the most comprehensive producer of meaning structures for several types of natural language available at present in terms of implementation, vocabulary, disambiguation of many-sensed words and references, etc. Moreover, as I have argued elsewhere (2), it is not the implementation of a conventional theory from linguistics, but is one with somewhat different principles of content.

In what I call its *basic mode*, the system already resolves anaphoras depending on superficial conceptual content of the text words. This is done in the course of setting up the initial representation. I shall call these type A anaphoras. For example, in "Give the bananas to the monkeys, although they are not ripe. They are very hungry", the system in its basic mode would decide that the first "they" refers to the bananas and the second to the monkeys. It can do that simply from what it knows about monkeys getting hungry because they are animals, and bananas having phases like ripeness because they are plants. All this information is, one might say colloquially, part of the superficial meaning of "banana" and "monkey".

This paper describes an "extended inference mode" of the system that tackles much more kinds of anaphora example than I shall call types B and C. Consider the correct attachment of "it" in "John drank the whisky from the glass, and it felt warm in his stomach". It is clear that the pronoun should be tied to "whisky" rather than "glass", but how it is to be done is not immediately obvious. Analysis of the example (see below) suggests that the solution requires, among other things, an inference equivalent to the sentence "whatever is in a part of X is in the X".

Anaphoras like the last I shall call type B, in that the inferences required to resolve them are analytic but not superficial. By "superficial" I simply mean that the sentence above, about parts and wholes, is logically true and not in any clear sense a fact about the world (or at least that rather about the meanings of words). What is meant by "superficial" in the distinction between types A and B will become clear after a discussion of the meaning formalism employed here.

We shall also discuss type C anaphoras, which require inferences that are no whitier at all, but weak generalisations often falsified in practice. I about the course of events in the world. Yet their employment here is not in any sense a probabilistic one. In "The dogs barked at the cat, and I heard one of them squeal with pain", we shall, in order to resolve the referent of "one" (which I take to be "cat" and "dog"), need a weak generalisation equivalent to "animate beings just like other animate beings may be unpleasantly affected". Such exponents are indeed suspiciously vague, and a reader who is uncertain at the point should ask himself how he would explain (say, to someone who did not know English well) the way he knew the meaning of "one" in that sentence. It can hardly be in virtue of a general rule that about cats and dogs because the same general rule applies to both. He saw whatever was chasing and being chased. I qualify the suggestion if he does not come up with something very like the inference suggested, and it may be the nature of natural language itself that is worrying him.

The inferences for type C, then, are general expressions of partial information (in McCarthy's phrase) and are considered to apply only if they are adequately confirmed by the context. What I mean by that will become clear in the course of what follows, but in no case do these expressions yield deductive consequences about the future course of the world. Indeed, they would be foolish if they did because the world's course cannot be captured in that way. In the unlikely example above, it might have been his earlier dinner that made him feel good.

BRIEF RECAP OF THE SYSTEM'S BASIC MODE OF ANALYSIS

In its basic mode, the system fragments texts (into phrase/clause-like items) and attaches a template to each. A template is a standard form of connectivity of semantic formulas as follows (where α denotes a complex item, to be described, corresponding to a chunk of an input word):

$$\begin{array}{c} F1 \alpha F2 \alpha F3 \\ / \quad \backslash \quad \backslash \quad | \quad \backslash \\ F11 F12 F13 \quad F21 \quad F31 \quad F32 \end{array}$$

F1, F2, F3 are the principal formulas of the template and are always agent, action and object (in that order), though any of them may be a formula or list of formulas dependent on main formula F1 etc. It should be said, in view of other current uses of "template", that it is not a surface structure at all, but a formal underlying meaning representation. Moreover, it does not function within a crude pattern matching scheme, such that if some text fragment has no templates matching it is thrown away, as it were. Special routines are called in such situations to construct an appropriate template item. All this basic material is set out in earlier papers.

The structure of formulas is explained below in some detail. In brief a formula is a nested list(a binary tree in fact) of semantic primitives called elements(expressed here as LISP atoms). Each such formula represents a sense of at least one an English word.

Let me give an example of a template structure at this point by using the following simplifying notation: any English words in square brackets [] stand for the meaning representation of those words in the PR-Centre Semantic system. This device is important in the interpretation of the material in that the content of the coded forms can be seen immediately, whereas the complex coded forms themselves would be difficult to read as [x], a sentence read a word at a time. But it is important to restate that the rules and formalisms expressed within () are really structured primitives, and that their tasks could not be carried out, as some still seem to believe, by massaging the language words themselves to stand for their own meaning representation.

So then, the template for "The black horse passed the winning post easily" could be written (ignoring any ambiguity problems for the moment):

```
[horse]----[passed]-----[post]  
          ↑           ↑           ↑  
[the black] [easily] [the winning]
```

In any or all of the agent, action or object formulas are missing, the template node(s) is filled in with a dummy element DTHIS. Thus a template for "The old house collapsed" could be written:

```
[house]----[collapsed]-----DTHIS  
          ↑  
[the old]
```

In the case of structures like prepositional phrases, we consider the preposition represented as a pseudo action, and the whole template as having a dummy agent. Thus for "at the Derby", we have:

OTH|S.....(a).....[Derby]

↑
[the]

The representation of a text (composed of fragments) is then a network of these template networks. The templates are interconnected by case ties. The notion of case is discussed in more detail below, but for the moment a case can be thought of as tying one template to some particular node in another template by a link of a certain type ; namely the case type, which specifies the sort of dependence the template has on the latter. In the sentence "He lost his wallet / in the subway" (fragmented at the stroke) we might say that the second fragment of the sentence depends on "lost" in the first, that the dependence is the locative case. Thus in the representation, the template for the second fragment would be tied to the object, affix, node of the first, by a link labelled LOCA. The node on the first template to which the case tie ties is called the marker, the second template.

Type A anaphoras are dealt with adequately within this framework, though it may be the purpose of this paper is not to describe the basic model or operation of the system because we get a denser network of links by considering the formula for the appropriate referent substitution than the problem pronoun than with its rivals. A link is preferred if it is listed. So, if we think of the formula for "ripe" in terms of its preference for application to plants, we see why a denser network arises in the above example for correct solution, rather than for one equivalent to "ripe monkeys". The way in which formula records preferences of various sorts is described below.

There is also, these type A anaphoras also constitute links between templates, from the pronoun variable to its correct referent. This is the final first form of the whole representation obtained from the basic input :

(CASE MARK ANAPHORA F1 F2 F3 (F1 dependents) (F2 dependents) (F3 dependents))

This initial form of representation of a text paragraph is called its IREP (Initial Representational Representation), or "semantic block". No emphasis has been placed here on syntax analysis of the input, and a reader (the compiler 1, 2, 3 or 4) will see that all of conventional grammar analysis has been done in the course of setting up this form of representation. An example of such an IREP, for the monkey-banana example, is given below as computer output.

Having sketched in the basic mode and its representation, we can now get to our buttons and sketch in the extended inference mode that is the heart of this paper.

BUCKLETHWAITE OF THE "EXTENDED INFERENCE MODE"

The extended inference procedure is called whenever the basic mode cannot resolve an anaphora between two or more candidates by semantic density. In the example about John and his stomach, density techniques may no longer decide whether the glass or the whisky is in his stomach. On a basis of preferred agents and objects of actions (what I have referred to as superficial conceptual information) both are equally good candidates. The extended inference procedure is called, and if it succeeds it returns a solution to the basic mode which then continues with its analysis. If it too should fail to reduce the number of candidates to one, then the top level of the system tries to solve the problem by default, or what a linguist would call focus. Roughly, that means to assume that whatever was last talked about is still being talked about. So, in "He put the bicycle in the shed and when he came back next week it was gone", neither agent criteria, nor the extended inferences to be described here, will help at all. So the system may as well assume, in this particular context, that the bicycle is still the focus of attention, and thus the reference of "it".

Consider again the following sentence after all the basic mode's constraints have been applied:

[1] John drink the whisky / 2 DIRE : DTHIS from a-glass / 3 :
[3] -> [3] IN 1: DTHIS in his-stomach]

Now, as in the square brackets, this item is a template for extraction. The case names DIRE (direction) and IN (containment) respectively. The numbers 1, 2, 3 refer to templates 2 on 1, and 4 on 3, respectively. The DTHIS's are dummies added to fill out the canonical structures for cases of missing agents, objects etc. Further assume that the "it" has been tied to "John" by the basic mode, and that no information of analysis. And assume too that the basic mode has decided that of "canonizes" for the references of "it" ("whisky" and "shed"). Of course if there had not been such a list of more than one word in the routine under description would not have been able to decide.

Extraction now then picks from each template in turn, and it and only it, unpacks. It obtains a representation of either an answer word or the whole pronoun itself. An extraction is the unpacking of the template case into both those in the action (second) formula and the template and those labelling as link to other templates. In our example we obtain the following extractions which are represented in forms as follows (where the first digit refers to the argument #, the second to the extraction, and "*" links words with a single argument):

1: drink-ki (IN in-) John+part]

2: drink-ki (DIRE to) John+part]

C1 (Inh₁(IN₁ from glass))

Alt 1.2 IN₁ in his stomach.

Now in this informal representation we have acquired new formula-like objects that express, in canonical form, new analytic information extracted from the existing templates, and from which new inferences can be made. It is postulated that the generation of this recipient information from the deeper levels of the formulas is essential to the process of understanding. These new forms differ from standard formulas only in that their second node, or "node of action", has had a case name CONS₂ onto whatever the node was originally. Note here that the form (IN₁ in) is not redundant since the IN indicates the sense precisely as containment, while the English preposition can indicate many cases other than containment, as in "in five minutes".

We cannot decide just how these particular extractions were made, even in the absence of full detailed knowledge of the structure of formulas:

If we furthermore wed from the template for "John drank the whisky" (assuming again the structure of the formula for "drink") it follows that the liquid drink is subsequently inside the drinker. This is because, when making up the formula for the action "drink", we express in it just the action consisting in causing a liquid to be inside the agent of the action.

Again, this is inferred because the same formula specifies that the liquid enters the drinker through a specific part of the drinker (his mouth, of course).

C1 is inferred from the direction case of the second template, whence we know that it was the whisky that was moved from the glass.

Lastly, Alt 1 is inferred from the direction case of the fourth template, because whatever the referent of it is, it is also in John's stomach.

Let us see where we are: we have obtained new template items that yield no extra information, but did not appear in the original text. (As we shall see in the detailed treatment below, some of the above are obtained from extraction, more strictly defined, and some from what I shall call "repacking the semantic block".)

In the pool for inference procedures we now have the original templates that mention either the variable pronoun or the possible "answer" referents, plus the extractions. We also have access to an inventory of Common Sense Inference Rules (CSIRs) which are of the form T (T₁ T₂), where T₁ and T₂ are T-forms, that is, templates or extractions.

We now try two strategies in turn: first we try a zero-point strategy, which is to try to identify an answer template (or extraction) and a variable template (or extraction) without the use of inference.

The general assumption here, and it is a strong psychological assumption, is that in order to resolve these painful ambiguities the inferential system is going to use the shortest possible chain of inferences at all times. And a zero-point strategy will, as it were, have no length at all (in terms of a chain of CSI inferences) and so if it succeeds, it will always provide the shortest chain.

This criterion is adequate (or the example under discussion, because it can hardly be given a definition of matching). Identify extractions as $\{X\}$, and identify $\{Y\}$ and the whisky, and we are home. This is the resolution of a B-type anaphora, requiring only analytic, non-inductive conceptual information. It should be noticed that this same ambiguity (defined earlier by the need of weak inductive information for their solution) can also be solved by the zero point strategy, because some extractions, and in particular those from the head rule (see below), are inductive and non-analytic.

If the zero-point strategy fails, we bring down all the CSI rules that contain in action subformula occurring in an answer or problem formula in the reply, and attempt to find the shortest chain that leads from some answer to some variable.

Thus, in the sentence "The soldiers/ fired+at the women/and I saw several fall", we extract a form equivalent to [soldiers strike women], since we can tell from the formula for "fired+at" that the action was intended to strike the object of the action. We are seeking for partial confirmation of the assertion [$?several$ fall DTHIS], and such a chain is completed by the rule [X strike Y] \rightarrow [Y fall DTHIS], though this is rule equivalent to, say, [X strike Y] \rightarrow [Y die DTHIS]. Since there is nothing in the sentence as given to partially confirm this rule in a chain, and cause it to fit here. Since we are in fact dealing with subformulas in the statement of the rules, rather than natural language words, "fitting" means an "adequate match of subformulas".

It is conceivable that there could be an, implausible, chain of rules and extractions giving the other result, namely that the soldiers fired: [$\{\text{soldiers fire DTHIS}\} \wedge [X \text{ fire DTHIS}] \rightarrow [Y \text{ fire+at } X] \rightarrow [X \text{ strike } Y] \rightarrow [X \text{ fall DTHIS}]$ etc., based on the assumption that "fire that fire guns yet fired at (...he who lives by the sword shall perish thereby...)". But such a chain would be longer than the one already constructed and would not be preferred.

MORE ON THE BASIC MODE

Formulas

Formulas are structures corresponding to senses of words, expressing their semantics. Much of the body of this paper is concerned with the manipulation of such structures, and the extraction of information from them, so it is important to have some general idea of their construction and interpretation.

Formulas are binary trees, expressed as lists, of semantic elements, parenthesized by right and left brackets. The elements are either case items, or actions, such as CAUSE, STRIK, CHANGE, or items such as THING, HAD, WRAP, Imitating, as examples here, element names that are often spelt out in Anglo-Saxon monosyllables, but there are about 1000 of them. There is one brief informal explanation, such as GRAIN, used to denote a grain. There are also elements like KINO indicating possession, and elements (indicated by an initial *) that stand for sets of elements such as *ANIMATE to cover MAN, BEAST and other human beings. In addition, most elements have a negated form, indicated by the element name. I assume here that the use of elements or primitives of this general sort, that are not logical connectives, need no formal defense at this point.

The most important element in a formula is its rightmost, called its head. This is also the most general sort of item, or action, or type of element. It expresses what corresponds to: for example, any word referring to a human being will have MAN as its head.

Since formulas are binary trees of unlimited depth, they can be decomposed into pairs of elements and subformulas, down to the most primitive elements. This process is called splitting up the formula, or decomposing it while reading it. At each stage there is a dependence of the left element on the corresponding right: this dependence is called a dependency.

Formulas depend directly on action

For example, MAN is an item, or not an item.

MAN is an item, or not an item or not an item.

In this last example the interpretation is unambiguous, since we know the range of functions of the elements in play. So, the formula (WRAP MAN) always means "a human" envelops something, and WRAP is always an action when in the right hand position, and WRAP a quantifier when in the left, dependent, position. WRAP is never an item and MAN is a agent not an object in the English. It would be an object if in the representation of "a human being is enveloped" because agents of actions may be unmarked, but objects are never unmarked. Conversely (WRAP THING) is a

container, since WRAP is always a qualifier when in the left hand, dependent, position in a pair.

An important notion is that of the semantic preferences that formulas can express. Consider the formula for "grasp" of objects:

"grasp"action₀ → ((*ANI SUBJ) ((*PHYSOB OBJE) ((THIS (MAN
WRAP) PHYS) (DOOR SENSE))))

In each element, SUBJ and OBJE occur at the top levels on the left of the formula, and at that level in an action formula they express the preferred agent and object of the action concerned.

The grasp action in this sense is an action preferably done by animate beings (*ANI) to physical objects (*PHYSOB), and consists in a act of holding, by taking, and done with an instrument (INST is the case element) which is a part of the body. When I say "prefers" here, I mean that, if the preferred agent or object cannot be found, the template is constructed with whatever is available. Thus, "The robot grasped the block" would never be rejected; it would only be less preferred than any possible competing interpretation that had an animate agent. I have argued in (2) that this approach to rules of formation has unexpected consequences for linguistic theory.

But one could also expect another formula to be available for "grasp", one such as:

"grasp"action₀ → ((*HUM SUBJ) ((STGN OBJE) (TRUE THINK)))

in this case, we have an action, preferably done by human beings to some object (e.g. ideas, principles etc), namely of thinking them through, or "grasping" them.

The grammars of the basic mode always fill this last formula into a template structure for "He grasped the principles", and the other template for "grasp" into the template for "the boy grasped the toy", by means of the preference and semantic density techniques described in chapter 5. These preferences for agent and object are part of the "superficial conceptual information" referred to earlier.

A few other rules will help to clarify the notion of "knowing our way around a formula" when interpreting its

agents. The implicit (need not be specified by SUBJ case) agents (if they occur at the top level in an action formula as explained above, or they attach to the head of a formula, as in:

"grasp"action₀ → ((NOT PLEASE FEEL) (SUBJ MAN))

then, the grammaticality of agents being to the left of (= dependent of) the corresponding action, is violated, since MAN is the agent for FEEL, while at the same time being the head of the whole formula.

This violation of order in search is indicated by also violating the ordering restriction that normally makes the SUBJ case element the dependent (right-hand member) of the pair in which it occurs. The corresponding rule of analysis is "On encountering SUBJ as dependent, expect action for the agent to follow to the left".

Moving objects are never implicit. Moreover, an object is considered the object of all actions to its right. This enables us to express the important notion of real and apparent agents of actions. For example (19):

"fire+at" (animate)

*((ANIMATE) (ANIMATE) (OBJECT) ((STRIK GOAL) ((THING MOVE) CAUSE)))

This action (the preferably by human beings to animate beings) is one of striking something (the bullet) with the aim (GOAL case) of striking something. Since *AN1 is the object of all actions to its right, it is the object not only of CAUSE, but also of STRIK. Hence, the bullet is shot by the same animate being. Moreover, THING(the bullet) is internally the agent of MOVE, not the object of CAUSE, which is correct as far as the meaning of "fire+at" is concerned.

At present we associate with a distinction system of ten cases, which are listed below, together with (in capital letters) the semantic elements that complement them, the questions that define them, and examples of subformulas expressing them. Defining a case is a tricky matter, and our present method is reasonably adequate. Note that the subformula examples are of those parts of a formula that would occur at that position AS PART OF THE MEANING OF A WORD. The subformulas do not, of course, show how the system would express the quoted words if encountered in a text, when they would be represented by a template.

recipients: FOR "for a woman" → ((FEM MAN)FOR)
what/whom to? what/who for?

instrument: INSTR "with a stick" → ((LINE THING)INST)
what with? in; what means?

direction: UP (see below), TO, FROM, UP
"from the top" → ((UP POINT)FROM)
down to where; from? at what? out of where? by what?

possessives: POSS "owned by a man" → ((MAL MAN)POSS)
who owns the thing mentioned?

locations, e.g. "at that time" \rightarrow ((THIS(WHEN POINT))LOCA)
where? where? when at? by what? in what time? near what? at what
time? during what? before when?

and similarly, (1) "in a glass" \rightarrow (((((FLOW STUFF)OBJE)WRAP)THING)IN)
in what?

and so on: (2)
out of what? for what?

(3) "to strike"
"to" is to strike a woman" \rightarrow
to what end? for what purpose?

and similarly, (4) "WITH"
"without a glass" \rightarrow (((((FLOW STUFF)OBJE)WRAP)THING)NOTW, TH:
unaccompanied by what/who? with what/whom? without what/whom?

and so on: (5) who did this?

etc.: (6) and what has this done to?

Certain elements above have negative forms leading to additional
elements NOTOR, NOTPOS, NOTIN, NOTWITH.

These elements have two functions, and occur in two sorts of
formulas: formulas and IREPs. In formulas they express part
of the meaning in a broad sense. Thus in

"move" \rightarrow ((MOVE TIE)INFL(FLOW STUFF)SOUR)THING)

"move" in a formula has a liquid source (FLOW STUFF), and is in a
directional template (TIE). The other function of these elements is, as
already explained, the name of the tie between the template for some
formula and the part of another template

and in the case of the class of direction case elements (TO and
FROM) and POSS only as the indicator of the case of a fragment,
never in formulas. Conversely POSS occurs only in formulas, never as
the indicator of a fragment case.

The last element is only included in a formula when it is specific:
the case element itself aspect of the case is involved. In the formula
for "move", for example, we include a direction specification for
movement destination POINT(TO). However, in the formula for "move" we do
not include the element TO or FROM, even though movement must in fact
be in some direction, since have no reasonable expectation about it
as to what "point". Sentences containing "move" may very well go
on to describe the direction involved, but its association with "move"
is completely arbitrary and we can't expect any confirmation of
expectations that would, say, resolve ambiguities. In this respect

the system differs from other systems that do create case expectations for aide classes of actions, which are essentially unhelpful, as in this example, and so we would claim unhelpful automatically.

Fig 180. An semantic block representation

Block 180 below is an example of an IREP for a pair of English sentences. The format of the block is the list structure described earlier - the result from the basic mode of operation. The only difference from that format is the presence in it of the stereotypes from which French is consequently generated (see Wilks and Herskovits (1977) and St. John (1978)). The French, as generated from the block, is written above the printout of the block itself for diagnostic purposes. The generated context-sensitive stereotypes are drawn into the block due to analysis, along with the formulas. The process of generation is through successive unwrapping of the block.

180. DONNER TO THE MONKEYS ALTHOUGH THEY ARE NOT RIPE! THEY ARE VERY HUNGRY.

DONNER TO THE MONKEYS AUX SINGES BIEN QU' ELLES NE SOIENT PAS MURES !
AL THOUGHT FEUILLANT.

(((SINGE~1 BANANAS) ((EX NIL NIL ((MPCL)))) 4 (((THIS DTHIS
BUTTERFLY~1 FEUILLANT (*ENT OBJE) GIVE)) GIVE (DONNER)~) ((MUCH ((*ANI
MONKEY~1 FEUILLANT WANT)) (OBJE PLANT)) BANANAS (FEMI BANANE)) NIL
((PRES (BE BE))) ((PTO ((GIVE) REC) ((&PREOB A))) 6
((CETTE (THIS) MURTHING) ((THIS P00) PTO NIL)) ((THE (MUCH ((MAN LIKE)
BEAST~1))) (OBJE~1 (MASC SINCE)) NIL NIL NIL)) ((ALTHOUGH THEY ARE
NOT RIPE~1 ((TASTE SENSE) WANT)) (OBJE PLANT)) ((THEY BANANAS)
((PRES (BE BE))) ((NPRES~1 (BE BE)) ARE ((IS_OBJECT HUNGRY)
AVOIR (D1ROB Q FAIM)) ((IS_OBJECT THIRSTY) AVOIR (D1ROB Q SOIF))
((IS_OBJECT AFRAID) AVOIR (D1ROB Q PEUR)) (ETRE)) ((PLANT POSS)
((CETTE (CAN USE) KIND) RIPE (MUR)) NIL NIL NIL)) ((THEY ARE~ VERY
HUNGRY~1 ((~ NIL NIL ((INCOL)))) 1 (((THE (MUCH ((MAN LIKE)
BEAST~1))) (PRON & MASC PLUR~)) ((PRES (BE BE)) ARE
((IS_OBJECT HUNGRY) AVOIR (D1ROB Q FAIM)) ((IS_OBJECT THIRSTY) AVOIR
(D1ROB Q SOIF)) ((IS_OBJECT AFRAID) AVOIR (D1ROB Q PEUR)) (ETRE))
((WANT POSS) ((TASTE SENSE) WANT) STATE) KIND)) HUNGRY (AFFAME~)
NIL ((CETTE (CAN HOLD) VERY (TRES))) NIL)))

THE IMPLEMENTATION OF THE EXTENDED MODE

There are three parts to the extended inference mode: the REPACK routine that takes the IREP block and repacks it; the EXTRACT routine which handles extractions, new knowledge not explicit in the text source; and INFER which tries to link an answer T-form to one or more question variable, that is, a text pronoun giving results.

The REPACK routine.

This attempts to replace dummy nodes in the IREP wherever possible before handing the whole representation to the extraction procedure. The replacement is itself a complex form of inference, sometimes as simple as the inference routines on which we are concentrating here. However, it is clear by inference that these procedures are mutually dependent and in a hierarchical fashion, so we may legitimately consider the construction in one single process in this way. The degree of "filling in" resulting done by REPACK, in the construction of a machine-level T-form from IREP, varies with particular action cases.

Let us look at the informal extractions done from the "John drank the whisky." example. We will see that the new T-forms 21 and 22 are actually obtained by filling in a dummy agent in some template from a node in another template. Thus from [2 DIRE : OTHRS from WHISKY] we obtain the new T-form, numbered 21, [whisky (DIRE from) WHISKY]. This was done by filling the dummy agent node of the template for "with a glass" with the formula for "whisky", and shifting the direction case marker into the pseudo-action. This is a packing, not an extraction proper, since the T-form obtained simply replaces a template "assert on" already in the representation. As we shall see, a true extraction is a new T-form altogether.

Let us now distinguish replaceable and unreplaceable cases roughly as follows. The dummy agent in the second (instrumental) template for "He bought father/with a club" cannot be replaced to yield any form equivalent to [father (INST with) aclub]. So we may say that the instrumental case is unreplaceable. But the dummy agent in the second (agent) template of "He bought the flowers / for his mother" can be replaced to yield a form for [flowers (RECIP for) his mother], and the agent case is replaceable, and is replaced by the operation of REPACK.

At the moment REPACK can be written in LISP as:

```
REPACK(IREP)→IREPR
```

3.2.2. The Extract routine

The routine takes each T-form, or template-like item, in IREPR in turn and processes it, modified if necessary, in a new block IREPE, followed by T-forms extracted from it. At present, extractions are made only from templates that contain either one of the possible 30 variables or one of the variables of the problem. The former are templates containing a formula for a word on the list ANS, the latter are templates one of whose nodes is (QUERYMARK THIS). Any templates that contain either an answer or a variable are simply transferred unaltered from IREPR to IREPE.

Thus, the general form of the extraction routine at the top level is

EXTRACT(IREPR ANS) → IREPE

Taking each template in turn, we first consider those processes that modify it, and then those which produce new T-forms from it. In the first case, there are some manipulations to do with negation, and with the SUBJ and OBJ cases.

If the agent or object formula is negated, the negative item in its formula is removed and the head of the corresponding action formula is inverted, so that all the subsequently applied inference rules can correctly negate the action. Thus, in \exists notation, we would achieve the following procedure: the coded equivalent of

(John drank not(gin)) → [John not+drink gin]

Each agent and object formula is then scrutinized by the question "does it satisfy the preference expressed by the corresponding action? If it does not, does any "of-phrase" qualifier of it do so instead? If so, replace the agent or object by that "of-phrase" instead as the true agent or object.

Then:

[John drink a glass of wine] → [John drank wine]

[Women drink wine] → [Women drank wine]

The main purpose of EXTRACT takes the action formula of a template and moves leftwards through it seeking case heads (other than SUBJ and OBJ). If it finds one, it asks is it replaceable, and, if it is, EXTRACT looks at subsequent fragments to see if REPACK has already replaced it. If it has been replaced it is forgotten, thus avoiding the same case information being extracted twice. It detects that REPACK has made such a replacement by finding the case name itself in the predecessor of a succeeding template, and a replaced dummy as the corresponding Agent.

With the next case, for example, the dependent of the case element becomes the action of the new T-form. In this case, as with every other, an attempt is made, on finding a potential agent or object for the new T-form at the top level of the action formula of the template formula, in position, to identify it with the main agent or object formula of the template. If this can be done, the agent or object formula of the original template is used, as being more specific. For example, in extracting from the action formula in the template for "John to shoot the deer", we find the goal case in [fire+at], with the term *it* STRIP, which is the attempted action. The object of that action, found to its left, is *ANI which can be instantiated by the formula for "deer" in the main template, namely (THIS BEAST). So the latter is used as the object of the new, extracted, T form [John trikes deer], since "deer" is more specific than "animate being".

For most other cases (agent, direction, location, containment and concepts), the case element provides the new pseudo-action, and the other object of it is specified as the dependent of the case element. The dependent is found as follows: it is the highest level object actually available in the action that dominates the case (to its rightmost part in the formula).

so in the formula for "pour" in "I pour the wine"

(I POUR (WRAP (FLOW STUFF))) ((((WRAP THING) TO) MOVE) CAUSE))

we consider moving leftwards through the the formula) the direction case in (the first move) (WRAP THING) TO), implying that the (FLOW STUFF), which is the highest level object in the formula, is moved in the direction of a container, or (WRAP THING). The case element TO is dominated by MOVE, whose highest level object is (FLOW STUFF) which would become the pseudo agent of the new extracted formula, but this it can be identified with the object of the old template, namely "wine", it is, and that becomes the pseudo-agent of the new T-form, since it is more specific than "liquid". It is, (WRAP THING), the container, becomes the object of the new T-form and the dependent case element becomes the action so we get an extracted formula

(CAUSE (WRAP (FOLLOW STUFF)) (TO P00) (WRAP THING)))

which is (I POUR (WRAP (some+container)))

The Rule Engine

This section has access to the representation IREPE produced by CEDAR, and CSIR, the inventory of common sense inference rules. Its domain at the top level is

INFER(IREPE ANS CSIR) -> ANS'

given. It's a longer ANS or some sublist of it, preferably containing minimal complete item, the solution.

THE zero-point tries the zero-point strategy trying to match some argument with some variable T-form directly, with no use of CSIR rules. This being here means that the two T-forms as arguments of a function try to produce a non-nil result, which will be a list of the corresponding (but non-identical), nodes in the two matched T-forms. The solution of the example "John drank the whisky. . . ." is illustrating the zero-point strategy, and rests on the matching of the two T-forms:

((whisky ((N in) John+part))

((it (N in) stomach))

the `(BASIC)` returns the list `((?it whisky ((John+part stomach)))` confirming the answer. No such match can be made for the alternative relation "glass".

The basic principle of inference at work is to select the shortest possible chain of inferences, on the assumption that an ambiguity of understanding of this sort should be solved in the most shallow way possible. If the situation becomes intolerable for the understanding, thus a zero-point solution, if available, will always be the shortest possible chain of inference.

If the zero-point strategy fails, the CSIRs are called, stored as a list consisting of their action subformulae, and, moreover accessible from inside "antecedent" and the "consequent" action subformula. At present we can make inferences of length one + those which are equivalent to a single CSIR for their solution. However, it should be possible to extend the present strategy to at least length two; and unfortunately they might almost never be any longer.

Let us look once more at the example "The soldiers fired at the women and I am several anti". We have to resolve "several", which cannot be done by the basic mode since both soldiers and women can equally well be 1. Let us set out the fragment representations and the resolution obtained as follows:

((fire ((soldier) (at) (woman)))

((I (am) (several) (anti) (THING))

((fire ((several) (at) (woman)))

((I (am) (several) (anti)))

((fire ((several) (at) (THING)))

((I (am) (several) (ANTI) (OTHIS)))

The inventory of rules is searched for those containing any action component occurring in a T-form in the pool that also contains either an "animate" or a "problem variable". In this case we pull in a rule internally expressed:

[1] STRIKE (*ANI 2) ~ ([*ANI 2] fall DTHIS)

Here variables are indicated by numbers ; *ANI expresses a restriction on the variable that any value of it must be animate, and the double ~ indicates that this rule can be considered as running in either direction.

This CSIR form is of course a more perspicuous form of:

[1] (THIS STRIKE (*ANI 2)) ~ ((*ANI 2) (NOTUP BE) DTHIS)

which would of course cover a wider class of activities than simply the English verb "strike". It would cover at least "hit", "batter" etc. etc.

But a chain of length one is established by the rule from T-forms 12 to 13 since the "animate condition" is satisfied and the variable ANIMATE is identified by the rule within the formula for "woman". It should be noted here that the inference rules are very weak in that the application of a rule like the present one is perfectly consistent with the description of a situation where an animate being is struck in some way but does not fall. And this weakness is wholly intentional.

One important inflection in CSIRs is whether or not negation is significant in them. The negation of the action in a T-form is not being significant. Consider "John drank no gin / in his martini / but it left him / in his stomach nonetheless". In the template for the first fragment, shifting the negation to the action, and extracting for the containment case from the formula for "drink", we will obtain a T-form

(gin (MARTINI BE) container)

and another :

(gin MARTINI BE) John]

Conversely no null action. By the same method, from the second template (not (A) (IN BE) container) and [martini (IN BE) John], in trying to tie up one of these drinks by matching to the extraction in ?it (IN (in) John), we shall, without the use of CSIRs be able to tell the ANS is down to a single member, namely [martini], since MATCH will show us that (gin) cannot stay on ANS.

However, it is not giving the same analyses to a sentence like "John
walked over to the window / and he knew / that he would get "it",
in which the complement of "it" is "car" and not "window" and he
would get it, by linking the first and fourth T-forms with some
extra relation (call it λ) such that $\lambda \rightarrow \{(*ANI\ 1)\ have\ 2\}$

But, just this is the point, if the same sentence had concluded "...
and he knew he would get it" we should have required the same
rule, and the same answer, this rule has its "consequent" action
on the whole sentence, but the question of it is irrelevant to its application.

These two examples should can be seen to be non-deductive very clearly
but they could also be said to be of the form $A \rightarrow \{B \vee \neg B\}$,
which is true with any content whatever in a deductive system,
but is not true and necessary.

3. THE MODE OF USE AND SOME PARTICULAR COMPARISONS

The problem of mode cannot be considered in any way adequately
here, since it because no one has any very clear idea of what
one means by the term. But even to qualify, the basic mode
must be shown to fail under a considerable vocabulary and range
of circumstances, and the extended mode must be shown to be
adequate to a decent sized inventory of CSIRs.

The standard errors that will also be those that will be expected to occur in course ambiguity procedure to back up the
basic mode, except ability to resolve ambiguity within a small
number of word fragments. Ambiguity over a context larger than
that is to be discarded, just as is ambiguity of the sort discussed
earlier, is not, but we should be prepared for it in an adequate
mode. That.

The standard errors of strong, and possibly naive, psychological
models, would that chain length is a reasonable metric to
choose between alternative interpretations. I think it is
reasonable, but that the tension introduced into understanding by
the need for economy of complexity has been overlooked. Notice here that
the length of a chain of CSIRs employed, not counting
extractions, different ways of writing down formulas will not
fit + chain length.

Indeed, I would justify the principle as being essentially an
extension of the so called semantic preference (Wilks 2) used in
the standard representation. That preference was justified as
leading to the "semantically deepest" interpretation which was, I
think, "the one with the least meaning" (in the sense in which a
word or word phrase carries the maximum possible information).
Indeed, the shortest chain of inferences also minimises the
extraneous inlay, and introduces the least extraneous inductive

information into the system. It is clear that such a notion of information-based choice is ultimately inadequate. We only have to consider a sentence like "I was named after my father" where it seems clear that we exclude one interpretation simply because it contains virtually no information. This alone shows there must be some qualification to a "minimising information" theory. However the fact that all available theories are wrong, by no means puts them all in the same position. I think such hypotheses about the overall manner in which an understanding system endeavors to maintain its coherence are well worth looking and testing, and that they represent an aspect of human language "competence" almost wholly ignored by current linguistics and artificial intelligence. One could make the point more precisely as follows: virtually all the systems in those areas define "success", that is to say the success of a particular parsed representation with respect to a text. What they do not tell us is what counts as a number of success are registered, as is almost always the case in realistic practice. But human understanders do not just pick the many, or opt for the first they find, or pick one at random - they prefer one in particular on some principled basis.

It is for this reason that the subject investigated in this paper must be treated in isolation from an adequate linguistic base theory, as I now seem to think. The inferring of a correct interpretation is intricately related to the systematic exclusion of competing interpretations, and any system that does not allow for the compatibility of sense and structure in at the start can hardly appreciate this point because the difficulty never arises there, but then neither does one essential aspect of natural language either. I have developed elsewhere (6) an abstract view of meaning along these lines: that to have meaning is essentially to have one meaning RATHER THAN ANOTHER. Or, put another way, having meaning essentially involves procedures for the exclusion of alternative interpretations. That, I believe, is the residual truth lurking beneath the "principle of non-contradiction", a thesis which when taken a face value is naturally false.

But we mention a closely related shortcoming of the micro-world approach to natural language analysis: it concerns what I believe to be an important motive in AI about the notion of "inference". Let me first of all state the obvious, not from dry motives of clarity, but because I believe the motive has important practical consequences in the area of natural language understanding.

There are often differences, in the bare sense of that word, of interpretation people might make from one assertion to another.

CECIL IS AN ENGLISHMAN. SO HE IS UNTRUSTWORTHY.

Cecil is an Englishman. SO he is untrustworthy.

This is triangular. SO it is three-sided.

To take it that all is a deduction, true in all possible worlds, and hence independent of the meanings of the words "Englishman", "Cecil" and "one another that".

That is an inference, simply and solely, and certainly not valid, whether or not it happens to be true for some English Cecil.

That is a valid inference, true in all possible worlds, as they say, because of the mutual agreement of "triangular" and "three sided", a fact that is captured by saying that the premise missing, for the triangle relation, namely "all triangular things are three-sided", is in effect always true.

What is the point for our purposes of all this dogmatic and semi-dogmatic classification? Simply this: the extra-conceptual elements in the so-called dictionary, that is) CSIR inferences of the sort we have discussed in this paper problems in text, are of type (iii). These inferences could function as part of a deductive system for the addition of sufficient inductively unreliable premisses to convert them to form (i). They could then function within established deductive machinery, such as first order PC, PLANNER in one of its modes of operation, etc.

But there may be no need for doing that, at least in the case of natural language analysis, because the conclusions reached can be no more reliable than the dubious generalisations functioning as premisses, whatever the power of the deductive machinery intervening.

In this paper I have described how such weak information can fulfil a problem-solving role in natural language analysis, in terms of a notion of "hypothesizing confirmed" inference in context. But that does not require the inductive machinery at all.

No point will be clarified here by noting two research situations under, by analogy, the deductive machinery may pay its way: (1) in robot and (2) in simulated micro-worlds.

In the case of a robot, really moving about in the world with possibly uncalculated information and plans, the world itself can provide a clear sense of contradiction. If the robot's deductions lead it the door is open, but it bangs into the firmly closed door in fact, then the conclusion is contradicted and the preceding premisses can be discredited, as would be the case with a scientific theory refuted in a carefully experiment. That is to say, the premisses may be uncalculated, but because there can be contradiction of conclusions the deductive machinery can transfer the "not" back to some premise, as in refuting a scientific theory, though the question of which particular bit should be transferred back to is very difficult of course.

the situation's maintain is quite different from the analysis of ordinary natural language where there is little or no expectation of contradiction; i.e. if, in understanding the text, the understander expects a contradiction. A. there is little or no chance of encountering contradictions in the near future. A robot could in principle learn contradictions by experiments, and in the case of a dialog in natural language, could in many step back and ask questions of the user, and accept and texts without experimenting on them. The problem is to ask one's informant all to discover such contradiction by a demand to go backtrack having done so, to the next least important information. And at the moment no system is in striking with this sort of capability.

The case of a simplified micro-worlds is different. Here there is no world to experiment, but there is no need for it since all premises are in effect known, and no real information can ever enter the system. Suppose, after executing the command "Clear off the top shelf of the book", it is clear by definition, apart from the possibility of a developmental contradiction. No lingering and sticky situations are to be found, to imperil the stability of the house of cards of the system. It will be clear that such situations have nothing to do with the unreliable inductive information required for the analysis of natural language.

It has already been seen that if there is no payoff to be obtained from a traditional approach to natural language understanding, there is no basis for pursuing it. This position is different from, but not inconsistent with, the position that distinguishes between one and another in formal fields, and urges the pursuit of the second, against the former: that is to say, the search for knowledge rather than work in a field rather than a search for an answer to questions about whose content we are ignorant. An argument in the negative position is that their method does also not depend on any sort of content, or human competence in this area, and not to be the formalisation of principles that could be expressed in a computer way.

After these classifications, some very brief comparisons follow between the ones described here, and three other AI approaches to language and finding: those of Charniak(7), Schank(8) and Hinde(9). In this comparison and criticism of systems is not intended to be, and I give only brief general remarks, in order to indicate the relations along a number of dimensions: (a) the degree of departure from the linguistic base constructed or proposed, in terms of application to everyday texts in English, (b) the degree of implementation and the definiteness of the task proposed as an approximation of the elusive notion of "understanding", (c) the location of the system within the inference-deduction opposition, and (d) the implementation of a preference system that both prefers certain interpretations to others on a reasoned principle.

All three schools do not consider the linguistic base essential, and is not particularly interested in the ambiguities of sense of words. The first school considers the most adequate linguistic base of the three, and the second school is the present system in general presuppositions. The third school is restricted to unambiguous simple words, and a detailed analysis of their meaning. Even if the words had only one meaning, it is doubtful if the meaning of complex concepts, such as culture, could be expressed in that way.

It is not intensively implemented but has a very detailed tool kit, the resolution of the sorts of anaphora which I mentioned in this paper. I think the strategies we advocate are not necessarily clear cut over whether or not the rules form an inference strategy as expressed at the word level, as he does, and whether or not he has the facility to chain CSIRs, or what he calls them. The system is on the verge of implementation through the work of a number of large programs. However, at present it is not a general specific task in the inference field. It is a class of inferences per se (as distinct from the drawing of conclusions of some problem or performance of some task) which is interesting. The "inferences to be drawn from x" is something which is not done, outside the pages of detective fiction, and is not completely implemented within its original context. As a separate and assessable task, its merits, as well as its disadvantages, in existing theory of grammar, is often said to be the same as in traditional syntax and the semantics (Minsky and Papert, 1969). I would remark, in an AI context, in that it is the traditional and traditional distinction (semantics and syntax) that such work in semantic analysis by computer (such as work for example) has found unnecessary in the last few years. It is interesting in view of the same authors' own interpretation-cognition distinction as well as the one I mentioned.

In this connection I do not feel on sure ground because I have not examined the authors' work, since, naturally their system is designed to answer this question of mine. The situation is complicated by the fact that some of the processes in Winograd's system cover processes that are almost certainly inference rules. My feeling is that Winograd's system is more general than Kripke's and Schank's, like the present one, are inference rules. The latter call in inference rules whose premisses are given to the machine context in hand. Any clarification on the relation of their work to this distinction would be welcome.

and it is not clear as of preference and choice between
the two. I think nothing has been done by the authors, and
I hope to prove that there will always be one and only one
coincident solution in terms of their rules, or that the first
coincident value, i.e. I think, the only worker in the field

and has given no attention to this question.

REF. 1

1. Meltzoff, J. An Artificial Intelligence approach to machine understanding. In: Schank, R., (Ed.), Computer Models of Thought and Language, San Francisco, 1971.
2. Hornbeam, D. A. Structure semantics. Stanford AI Laboratory Memo #161, 1971.
3. Hornbeam, D. A. Formal Semantics of Natural Language, Cambridge, 1973.
4. Hornbeam, D. A. Understanding without proofs, in Proc. Third Internat. Conf. on AI, 1973.
5. Hornbeam, D. A. and Herskovits, A. An intelligent analuser and generator for natural language, in Proc. Internat. Conf. on Computational Linguistics, Pisa, 1973.
6. Herskovits, A. The generation of French from a semantic representation. Stanford AI Laboratory Memo #212.
7. Hornbeam, D. A. Inability and natural language. Mind, 1971.
8. Hornbeam, D. A. Jack and Janet in search of a theory of knowledge. 1971.
9. Hornbeam, D. A. and Riegel, C. Inference and the computer generated sentence. In: Hornbeam, D. A. (Ed.), Computer Models of Thought and Language, Stanford AI Laboratory Memo #197.
10. Hornbeam, D. A. Generating natural language. Edinburgh, 1972.
11. Hornbeam, D. A. and Robert, S. Artificial Intelligence Progress Report No. 1, 1972. MIT AI Laboratory.
12. Hornbeam, D. A. Semantic memory. In: Minsky, M. (Ed.) Semantic Information Processing, MIT, 1968.